

EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	78	(histogram near bucket\$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 15:44
S5	269	(histogram near (bucket\$1 or bin\$1 or stor\$3)) and (bias)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 15:45
S56	9	((bucket\$1 or bin\$1 or interval\$1) with (full or fil\$2)) and ((bucket\$1 or bin\$1 or interval\$1) with empty) same (exceed\$3 with threshold)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 16:38
S72	25	(value\$1 with exceed\$1 with threshold) and (average with height) and (histogram\$1) and (bucket\$1 or bin\$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 15:53
S76	10	(minimum with percentage with row\$1) and adjusted	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 17:11
S77	101	(histogram\$1 near bucket\$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 17:05
S78	2808	(histogram near (bucket\$1 or bin\$1 or stor\$3))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 15:45
S79	19	((histogram near (bucket\$1 or bin\$1 or stor\$3)) same bias)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 15:48

EAST Search History

S80	0	(histogram near (bucket\$1 or bin\$1 or stor\$3)) and (high with bais)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 15:48
S81	42	(histogram near (bucket\$1 or bin\$1 or stor\$3)) and (high with bias)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 15:49
S82	17	(histogram near (bucket\$1 or bin\$1 or stor\$3)) and (high with bias) and percentage\$1	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 15:49
S83	19	(histogram\$1 with value\$1) and (height with row\$1) and (percent\$3 with row\$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 17:10
S84	40	(value\$1 with exceed\$1 with threshold) and (average with height) and (histogram\$1) and (bucket\$1 or bin\$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 15:53
S85	11	((bucket\$1 or bin\$1 or interval\$1) with (full or fil\$2)) and ((bucket\$1 or bin\$1 or interval\$1) with empty) same (exceed\$3 with threshold)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 16:38
S86	10245	((707/100,101,102) or (707/1,2)). CCLS.	USPAT; USOCR	OR	OFF	2008/01/29 17:09
S87	18	S86 and S77	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 17:10
S88	56	S86 and S78	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 17:10

EAST Search History

S89	0	S86 and S81	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 17:10
S90	3	S86 and S83	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 17:10
S91	65	(percentage with row\$1) and bucket\$1	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 17:11
S92	65	(percentage\$1 with row\$1) and bucket\$1	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 17:11
S93	13	(percentage\$1 with row\$1) and bucket\$1 and histogram\$1	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2008/01/29 17:12

[Advanced Patent Search](#)
[Google Patent Search](#)

Patents

Patents 1 - 3 on **high bias histogram buckets**. (0.03 seconds)[Sort by relevance](#) | [Sort by date \(new first\)](#) | [Sort by date \(old first\)](#)

Method and apparatus for identifying objects depicted in a videostream

US Pat. 6266442 - Filed Oct 23, 1998 - Facet Technology Corp.

Thus a potential source of **bias**; namely, prior pixel values from the ... Using a compressed **histogram** of the color of the face of a sign allows in a highly ...

Method and apparatus for identifying objects depicted in a videostream

US Pat. 6625315 - Filed Sep 16, 2002 - Facet Technology Corp.

Thus a potential source of **bias**; namely, prior pixel values 2o from the ...Using a compressed **histogram** of the color of the face of a sign allows in a ...

Method and apparatus for rapidly determining whether a digitized image frame ...

US Pat. 6449384 - Filed Jan 29, 2001 - Facet Technology Corp.

Thus a potential source of **bias**; namely, prior pixel values from the ... Using a compressed **histogram** of the color of the face of a sign allows in a highly ...[Google Patent Search Help](#) | [Advanced Patent Search](#)[Google Home](#) - [About Google](#) - [About Google Patent Search](#)

©2007 Google



histogram buckets and minimum rows

Search Patents

[Advanced Patent Search](#)
[Google Patent Search](#)

The "AND" operator is unnecessary -- we include all search terms by default. [\[details\]](#)

Patents

Patents **1 - 5** on **histogram buckets and minimum rows**. (0.02 seconds)

Sort by relevance | [Sort by date \(new first\)](#) | [Sort by date \(old first\)](#)

Method and system for sample size determination for database optimizers

US Pat. 6732085 - Filed May 31, 2001 - Oracle International Corporation

At step 304, The following is pseudocode for the top level routine for values with respect to the expected **histogram buckets**. In an 45 should be collected: ...

Query planning using a maxdiff histogram

US Pat. 6714938 - Filed Aug 31, 2000 - Microsoft Corporation

If the **histogram** is full, then the system proceeds to determine which of two neighboring **buckets** have the **minimum** variance (block 416). The two **buckets** that ...

Method and system for histogram determination in a database

US Pat. 6691099 - Filed May 31, 2001 - Oracle International Corporation

Each row will have the **minimum** and maximum value for that bucket, ... The number of **buckets** in the **histogram** could be based on the sample size. ...

Self-tuning histogram and database modeling

US Pat. 6460045 - Filed Mar 15, 1999 - Microsoft Corporation

The records of a table are also referred to as **rows** or tuples, The **B buckets** of the initial **histogram** are evenly spaced between min and max. ...

Method for superimposing statistical information on tabular data

US Pat. 7256784 - Filed Apr 26, 2004 - Hyperformix, Inc.

The effect of combining **histogram** data with standard statistical measures provides a ... Next, step 365 determines if the process is done with the **buckets**. ...

histogram buckets and minimum rows

Search Patents

[Google Patent Search Help](#) | [Advanced Patent Search](#)

[Google Home](#) - [About Google](#) - [About Google Patent Search](#)

©2007 Google



histogram buckets and percentage

Search Patents

Advanced Patent Search
Google Patent Search

The **"AND"** operator is unnecessary -- we include all search terms by default. [\[details\]](#)

Patents

Patents 1 - 10 on histogram buckets and percentage. (0.02 seconds)

Sort by relevance | **Sort by date (new first)** | **Sort by date (old first)**

Method and system for storing and reporting network performance metrics ...

US Pat. 6975963 - Filed Sep 24, 2003 - McData Corporation

Then, in step 230, each data point would be placed in one of these **histogram buckets**, and this compresses data as a single numeric value (eg, a **percentage ...**

Histogram-based approximation of set-valued query-answers

US Pat. 6507840 - Filed Dec 21, 1999 - Lucent Technologies Inc.

15 and 16 respectively show not significantly considered when **buckets** are ...

16, the **percentage** relative error for enough. For medium values of skew, ...

Real-time data cache size adjustment in a server computer

US Pat. 5732240 - Filed Nov 27, 1996 - Digital Equipment Corporation

Each **histogram** contains the frequency of occurrence of its respective operation categorized by byte-range **buckets**. Referring to FIG. ...

Data base optimizer using most frequency values statistics

US Pat. 4956774 - Filed Sep 2, 1988 - International Business Machines Corporation

Boundary key values may appear in adjacent **buckets**. Like the equi-width **histogram** method, the equi-height method assumes a uniform distribution ...

Multi-dimensional selectivity estimation method using compressed **histogram** ...

US Pat. 6311181 - Filed Mar 19, 1999 - Korea Advanced Institute of Science and Technology

Sampled DCT coefficients are computed for the **histogram buckets** count values with ... A **percentage** error is used for the accuracy of an estimation result, ...

Image query system and method

US Pat. 5579471 - Filed Mar 23, 1994 - International Business Machines Corporation

Within each block, the reduced bucket **histogram**, h , (say, 256 **buckets**) is computed. Given a query image or object, it is also partitioned into the same ...

Method of building multidimensional workload-aware histograms

US Pat. 7007039 - Filed Jun 14, 2001 - Microsoft Corporation

Progressively coarser grids are constructed over the data set and the densest cells are converted into **buckets** of the **histogram**. A certain **percentage** of ...

Method for superimposing statistical information on tabular data

US Pat. 7256784 - Filed Apr 26, 2004 - Hyperformix, Inc.

One such example is as a **histogram**. In viewing information and ... of **buckets** and reporting the number or **percentage** of samples that fall into each bucket. ...

Method and system for histogram determination in a database

US Pat. 6691099 - Filed May 31, 2001 - Oracle International Corporation

In one embodiment, each **histogram** uses a separate query, and therefore ...
upon the expected number of distinct values and the requested number of **buckets**. ...

Apparatus for reproducing an information signal stored on a storage medium

US Pat. 6957387 - Filed Sep 5, 2001 - Koninklijke Philips Electronics N.V.

A colour 15 **histogram** $H(l)$ is a vector $\langle H_1; H_2, \dots, H_n \rangle$ in which each element H_i contains the **percentage** of pixels of the colour C_i in the visual item ...

Google 

Result Page: 1 2 **Next**

histogram buckets and percentage

Search Patents

[Google Patent Search Help](#) | [Advanced Patent Search](#)

[Google Home](#) - [About Google](#) - [About Google Patent Search](#)

©2007 Google



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide

SEARCH
THE ACM DIGITAL LIBRARY
[Feedback](#)

histogram buckets

Terms used: **histogram buckets**

Found **349** of **238,273**

Sort results
by

☒ [Save results to a Binder](#)

Refine these results with [Advanced Search](#)

Display
results

☐ Open results in a new
window

Try this search in [The ACM Guide](#)

Results 1 - 20 of 349

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#) [>>](#)

1 [Approximation and streaming algorithms for histogram construction problems](#)



Sudipto Guha, Nick Koudas, Kyuseok Shim

March 2006 **ACM Transactions on Database Systems (TODS)**, Volume 31

Issue 1

Publisher: ACM

Full text available: [pdf\(1.38 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#),
[index terms](#)

Histograms and related synopsis structures are popular techniques for approximating data distributions. These have been successful in query optimization and a variety of applications, including approximate querying, similarity searching, and data mining, ...

Keywords: Data Streams, approximation algorithm, histograms

2 [Rangesum histograms](#)

S. Muthukrishnan, Martin Strauss

January 2003 **SODA '03: Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms**

Publisher: Society for Industrial and Applied Mathematics

Full text available: [pdf\(1.10 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#),
[cited by](#), [index terms](#)

A *rangesum query* to an array **A** is a pair (*l*, *r*) of range endpoints, which should be answered by $\sum_{l \leq i \leq r} A[i]$. To compress **A**, we consider representing an array **A** lossily ...

3 [Fast incremental maintenance of approximate histograms](#)



Phillip B. Gibbons, Yossi Matias, Viswanath Poosala

September 2002 **ACM Transactions on Database Systems (TODS)**,

Volume 27 Issue 3

Publisher: ACM

Full text available: [pdf\(472.82 KB\)](#) Additional Information: [full citation](#), [abstract](#),
[references](#), [cited by](#), [index terms](#)

Many commercial database systems maintain histograms to summarize the contents of large relations and permit efficient estimation of query result sizes for use in query optimizers. Delaying the propagation of

Ads by Google

[GIS Image Segmentation](#)

Shapefiles from satellite imagery
Wizard to segment, classify, batch
ImageSeg.com

[Imaging Analysis Software](#)

Get your complete and cost effective microscopy software.
[www.zechsoft.com](#)
[Document Scanning Service](#)

Free Online Quote.
Scan to PDF/TIF
Serving the DC Metropolitan Area
[www.ignitedscanning.com](#)
[Union Arbitrations Online](#)

Web-retrieval services provide fast retrieval for legal documents
[www.microsearch.net](#)

database updates to the histogram often introduces ...

Keywords: Approximation, histograms, incremental maintenance, query optimization, sampling

4 A new similarity measure for histograms applied to content-based retrieval of medical images



Joaquim C. Felipe, Agma J. M. Traina, Caetano Jr Traina

April 2006 **SAC '06**: Proceedings of the 2006 ACM symposium on Applied computing

Publisher: ACM

Full text available: pdf(97.72 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper presents a new similarity measure to compare gray-level histograms, aiming at reducing both false positive and false negative results, in the context of medical images.

Keywords: computer-aided diagnosis, content-based image retrieval, gray level histogram, similarity measure

5 Multi-dimensional selectivity estimation using compressed histogram information



Ju-Hong Lee, Deok-Hwan Kim, Chin-Wan Chung

June 1999 **SIGMOD '99: ACM SIGMOD Record**, Volume 28 Issue 2

Publisher: ACM

Full text available: pdf(1.18 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

The database query optimizer requires the estimation of the query selectivity to find the most efficient access plan. For queries referencing multiple attributes from the same relation, we need a multi-dimensional selectivity estimation technique when ...

6 Hierarchical binary histograms for summarizing multi-dimensional data



F. Furfaro, G. M. Mazzeo, D. Saccà, C. Sirangelo

March 2005 **SAC '05**: Proceedings of the 2005 ACM symposium on Applied computing

Publisher: ACM

Full text available: pdf(366.99 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

The need to compress data into synopses of summarized information often arises in many application scenarios, where the aim is to retrieve aggregate data efficiently, possibly trading off the computational efficiency with the accuracy of the estimation. ...

Keywords: histograms, multi-dimensional data, range queries

7 Histograms revisited: when are histograms the best approximation method for aggregates over joins?




Alin Dobra

June 2005 **PODS '05**: Proceedings of the twenty-fourth ACM SIGMOD-

SIGACT-SIGART symposium on Principles of database systems

Publisher: ACM

Full text available:  pdf(198.74 KB) Additional Information: [full citation](#), [abstract](#),
[references](#), [cited by](#)

The traditional statistical assumption for interpreting histograms and justifying approximate query processing methods based on them is that all elements in a bucket have the same frequency -- the so called uniform distribution assumption. In this paper ...

8 An information theoretic histogram for single dimensional selectivity estimation



Chris Giannella, Bassem Sayrafi

March 2005 **SAC '05: Proceedings of the 2005 ACM symposium on Applied computing**

Publisher: ACM

Full text available:  pdf(121.63 KB) Additional Information: [full citation](#), [abstract](#),
[references](#), [index terms](#)

We study the problem of one dimensional selectivity estimation in relational databases. We introduce a new type of histogram based on information theory. We compare our histogram against a large number of other techniques and on a wide array of datasets. ...


Keywords: entropy, histograms, selectivity estimation

9 Selectivity estimation by batch-query based histogram and parametric method

Jizhou Luo, Xiaofang Zhou, Yu Zhang, Heng Tao Shen, Jianzhong Li

March 2007 **ADC '07: Proceedings of the eighteenth conference on Australasian database - Volume 63**, Volume 63

Publisher: Australian Computer Society, Inc.

Full text available:  pdf(290.91 KB) Additional Information: [full citation](#), [abstract](#),
[references](#)

Histograms are used extensively for selectivity estimation and approximate query processing. Workload-aware dynamic histograms can self-tune itself based on query feedback without scanning or sampling the underlying datasets in a systematic and comprehensive ...


10 Optimal histograms for hierarchical range queries (extended abstract)



Nick Koudas, S. Muthukrishnan, Divesh Srivastava

May 2000 **PODS '00: Proceedings of the nineteenth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems**

Publisher: ACM

Full text available:  pdf(224.88 KB) Additional Information: [full citation](#), [references](#), [cited by](#), [index terms](#)


11 Data-streams and histograms



Sudipto Guha, Nick Koudas, Kyuseok Shim

July 2001 **STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing**

Publisher: ACM

Full text available:  pdf(147.56 KB) Additional Information: [full citation](#), [abstract](#),
[references](#), [cited by](#), [index terms](#)


Histograms have been used widely to capture data distribution, to represent the data by a small number of step functions. Dynamic programming algorithms which provide optimal construction of these histograms exist, albeit running in quadratic time and ...

12 Multi-resolution algorithms for building spatial histograms

Qing Liu, Yidong Yuan, Xuemin Lin

January 2003 **ADC '03: Proceedings of the 14th Australasian database conference - Volume 17**, Volume 17

Publisher: Australian Computer Society, Inc.

Full text available:  pdf(213.09 KB) Additional Information: [full citation](#), [abstract](#),
[references](#), [index terms](#)

Selectivity estimation of queries not only provides useful information to the query processing optimization but also may give users a preview of processing results. In this paper, we investigate the problem of selectivity estimation in the context of ...

Keywords: approximation query processing, histograms, selectivity estimation, spatial databases


13 Self-tuning histograms: building histograms without looking at data



Ashraf Aboulnaga, Surajit Chaudhuri

June 1999 **SIGMOD '99: Proceedings of the 1999 ACM SIGMOD international conference on Management of data**

Publisher: ACM

Full text available:  pdf(1.67 MB) Additional Information: [full citation](#), [abstract](#), [references](#),
[cited by](#), [index terms](#)

In this paper, we introduce self-tuning histograms. Although similar in structure to traditional histograms, these histograms infer data distributions not by examining the data or a sample thereof, but by using feedback from the query execution engine ...

14 Fast, small-space algorithms for approximate histogram maintenance




Anna C. Gilbert, Sudipto Guha, Piotr Indyk, Yannis Kotidis, S.

Muthukrishnan, Martin J. Strauss


May 2002 **STOC '02: Proceedings of the thirty-fourth annual ACM symposium on Theory of computing**


Publisher: ACM

Full text available:  pdf(266.50 KB) Additional Information: [full citation](#), [abstract](#),
[references](#), [cited by](#), [index terms](#)

(MATH) A vector \mathbf{A} of length N is defined implicitly, via a stream of updates of the form "add 5 to \mathbf{A}_3 ." We give a *sketching* algorithm, that constructs a small *sketch* from the stream of updates, and a *reconstruction* ...

15 Efficient histogram generation using scattering on GPUs

-  Thorsten Scheuermann, Justin Hensley
April 2007 **I3D '07**: Proceedings of the 2007 symposium on Interactive 3D graphics and games
Publisher: ACM


Full text available:  [pdf\(387.67 KB\)](#) Additional Information: [full citation](#), [abstract](#),
[references](#), [index terms](#)

We present an efficient algorithm to compute image histograms entirely on the GPU. Unlike previous implementations that use a gather approach, we take advantage of scattering data through the vertex shader and of high-precision blending available on ...

Keywords: GPGPU, histogram, image processing, real-time rendering, tone mapping

16 Global optimization of histograms


-  H. V. Jagadish, Hui Jin, Beng Chin Ooi, Kian-Lee Tan
June 2001 **SIGMOD '01: ACM SIGMOD Record**, Volume 30 Issue 2
Publisher: ACM

Full text available:  [pdf\(310.14 KB\)](#) Additional Information: [full citation](#), [abstract](#),
[references](#), [cited by](#), [index terms](#)

Histograms are frequently used to represent the distribution of data values in an attribute of a relation. Most previous work has focused on identifying the optimal histogram (given a limited number of buckets) for a single attribute *independent of* ...

17 A new histogram method for sparse attributes: the averaged rectangular attribute cardinality map

- B. John Oommen, Jing Chen
September 2003 **ISICT '03**: Proceedings of the 1st international symposium on Information and communication technologies
Publisher: Trinity College Dublin


Full text available:  [pdf\(153.70 KB\)](#) Additional Information: [full citation](#), [abstract](#),
[references](#)

Most current Database Management Systems (DBMS) use histograms in their query optimization, and in approximating query result sizes. This is because they can be utilized in determining efficient query evaluation plans. All the existing methods perform ...

Keywords: attribute cardinality maps, histogram-based algorithms, query optimization, sparse attributes

18 Multi-dimensional selectivity estimation using compressed histogram information

-  Ju-Hong Lee, Deok-Hwan Kim, Chin-Wan Chung
June 1999 **SIGMOD '99**: Proceedings of the 1999 ACM SIGMOD international conference on Management of data
Publisher: ACM

Full text available:  [pdf\(1.18 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#),
[cited by](#), [index terms](#)

The database query optimizer requires the estimation of the query selectivity to find the most efficient access plan. For queries referencing multiple attributes from the same relation, we need a multi-dimensional

selectivity estimation technique when ...

19 Mini-buckets: A general scheme for bounded inference



Rina Dechter, Irina Rish

March 2003 **Journal of the ACM (JACM)**, Volume 50 Issue 2

Publisher: ACM

Full text available: pdf(902.27 KB) Additional Information: [full citation](#), [abstract](#),
[references](#), [cited by](#), [index terms](#)

This article presents a class of approximation algorithms that extend the idea of bounded-complexity inference, inspired by successful constraint propagation algorithms, to probabilistic inference and combinatorial optimization. The idea is to bound ...

Keywords: Accuracy/complexity trade-off, Bayesian networks, approximation algorithms, combinatorial optimization, probabilistic inference.

20 Compact histograms for hierarchical identifiers

Frederick Reiss, Minos Garofalakis, Joseph M. Hellerstein

September 2006 **VLDB '06: Proceedings of the 32nd international conference on Very large data bases**

Publisher: VLDB Endowment

Full text available: pdf(1.22 MB) Additional Information: [full citation](#), [abstract](#), [references](#),
[index terms](#)

Distributed monitoring applications often involve streams of unique identifiers (UIDs) such as IP addresses or RFID tag IDs. An important class of query for such applications involves partitioning the UIDs into groups using a large lookup table; the ...

Results 1 - 20 of 349

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#) [>>](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2008 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

THE ACM DIGITAL LIBRARY

[Feedback](#)

histogram buckets and high bias

Terms used: **histogram buckets** **high bias**

Found 50 of 238,273

Sort results by

☒ [Save results to a Binder](#)

 Refine these results with [Advanced Search](#)

Display results

☐ [Open results in a new window](#)

 Try this search in [The ACM Guide](#)

Results 1 - 20 of 50

Result page: [1](#) [2](#) [3](#) [next](#) [>>](#)

1 [Histograms revisited: when are histograms the best approximation method for aggregates over joins?](#)

Alin Dobra

 June 2005 **PODS '05: Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems**

Publisher: ACM

 Full text available: pdf(198.74 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#)

The traditional statistical assumption for interpreting histograms and justifying approximate query processing methods based on them is that all elements in a bucket have the same frequency -- the so called uniform distribution assumption. In this paper ...

Ads by Google

[GIS Image Segmentation](#)
Shapefiles from satellite imagery
Wizard to segment, classify, batch
ImageSeg.com

[Imaging Analysis Software](#)

Get your complete and cost effective microscopy software.
www.zechsoft.com

2 [Cloud control with distributed rate limiting](#)

Barath Raghavan, Kashi Vishwanath, Sriram Ramabhadran, Kenneth

Yocum, Alex C. Snoeren

 October 2007 **SIGCOMM '07: ACM SIGCOMM Computer Communication Review**, Volume 37 Issue 4

Publisher: ACM

 Full text available: pdf(1.25 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Today's cloud-based services integrate globally distributed resources into seamless computing platforms. Provisioning and accounting for the resource usage of these Internet-scale applications presents a challenging technical problem. This paper presents ...

Keywords: CDN, rate limiting, token bucket

[Document Scanning Service](#)

Free Online Quote.
Scan to PDF/TIF
Serving the DC Metropolitan Area
www.ignitedscanning.com

[Union Arbitrations Online](#)

Web-retrieval services provide fast retrieval for legal documents
www.microsearch.net

3 [Synopsis for query optimization: A space-complexity perspective](#)

Raghav Kaushik, Jeffrey F. Naughton, Raghu Ramakrishnan, Venkatesan T. Chakravarthy

 December 2005 **ACM Transactions on Database Systems (TODS)**,

Volume 30 Issue 4

Publisher: ACM

 Full text available: pdf(530.68 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Database systems use precomputed synopses of data to estimate the

cost of alternative plans during query optimization. A number of alternative synopsis structures have been proposed, but histograms are by far the most commonly used. While histograms ...

Keywords: Cardinality estimation, histograms, sampling

4 RHist: adaptive summarization over continuous data streams



Lin Qiao, Divyakant Agrawal, Amr El Abbadi

November 2002 **CIKM '02**: Proceedings of the eleventh international conference on Information and knowledge management

Publisher: ACM

Full text available: pdf(345.75 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

Maintaining approximate aggregates and summaries over data streams is crucial to handle the OLAP query workload that arises in applications, such as network monitoring and telecommunications. Furthermore, since the entire data is not available at all ...

Keywords: adaptive approximation, data stream, histogram

5 Cloud control with distributed rate limiting



Barath Raghavan, Kashi Vishwanath, Sriram Ramabhadran, Kenneth Yocum, Alex C. Snoeren

August 2007 **SIGCOMM '07**: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications

Publisher: ACM

Full text available: pdf(1.25 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Today's cloud-based services integrate globally distributed resources into seamless computing platforms. Provisioning and accounting for the resource usage of these Internet-scale applications presents a challenging technical problem. This paper presents ...

Keywords: CDN, rate limiting, token bucket

6 On the hardware-software partitioning problem: System modeling and partitioning techniques



Marisa López-Vallejo, Juan Carlos López

July 2003 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 8 Issue 3

Publisher: ACM

Full text available: pdf(409.78 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

This paper presents an in-depth study of several system partitioning procedures. It is based on the appropriate formulation of a general system model, being therefore independent of either the particular co-design problem or the specific partitioning ...

Keywords: Hardware-software co-design, clustering, cost functions, expert systems, fuzzy logic, general optimization procedures, hardware-software partitioning, system modeling

7 Image-driven simplification



Peter Lindstrom, Greg Turk

July 2000 **ACM Transactions on Graphics (TOG)**, Volume 19 Issue 3

Publisher: ACM

Full text available: pdf(1.98 MB)

Additional Information: [full citation](#), [abstract](#), [references](#),
[cited by](#), [index terms](#)

We introduce the notion of image-driven simplification, a framework that uses images to decide which portions of a model to simplify. This is a departure from approaches that make polygonal simplification decisions based on geometry. ...

Keywords: image metrics, level-of-detail, polygonal simplification, visual perception

8 Accelerated warmup for sampled microarchitecture simulation



John W. Haskins, Jr., Kevin Skadron

March 2005 **ACM Transactions on Architecture and Code Optimization (TACO)**, Volume 2 Issue 1

Publisher: ACM

Full text available: pdf(214.93 KB)

Additional Information: [full citation](#), [abstract](#),
[references](#), [index terms](#)

To reduce the cost of cycle-accurate software simulation of microarchitectures, many researchers use statistical sampling: by simulating only a small, representative subset of the end-to-end dynamic instruction stream in cycle-accurate detail, simulation ...

Keywords: Reuse latency, sampled simulation, warmup

9 Using the visual differences predictor to improve performance of progressive global illumination computation



Valdimir Volevich, Karol Myszkowski, Andrei Khodolev, Edward A. Kopylov

April 2000 **ACM Transactions on Graphics (TOG)**, Volume 19 Issue 2

Publisher: ACM

Full text available: pdf(1.87 MB)

Additional Information: [full citation](#), [abstract](#), [references](#),
[cited by](#), [index terms](#)

A novel view-independent technique for progressive global illumination computing that uses prediction of visible differences to improve both efficiency and effectiveness of physically-sound lighting solutions has been developed. The technique is a mixture ...

Keywords: Monte Carlo photon tracing, adaptive mesh subdivision, density estimation, human perception, progressive refinement, view-independent solutions

10 Scalable near identical image and shot detection



Ondřej Chum, James Philbin, Michael Isard, Andrew Zisserman

July 2007 **CIVR '07: Proceedings of the 6th ACM international conference on Image and video retrieval**

Publisher: ACM

Full text available: pdf(2.68 MB)

Additional Information: [full citation](#), [abstract](#), [references](#),
[index terms](#)

This paper proposes and compares two novel schemes for near duplicate image and video-shot detection. The first approach is based on global

hierarchical colour histograms, using Locality Sensitive Hashing for fast retrieval. The second approach uses ...

Keywords: LSH, large image databases, min hash, near duplicate detection

11 Proceedings of the 9th annual conference on Genetic and evolutionary computation



Hod Lipson
July 2007 proceeding

Publisher: ACM

Additional Information: [full citation](#), [appendices and supplements](#), [abstract](#).

These proceedings contain the papers presented at the *9th Annual Genetic and Evolutionary Computation Conference (GECCO-2007)*, held in London, UK, July 7-11, 2007. For the first time GECCO was held outside the US. This clearly proved ...

12 Gaussian random number generators



David B. Thomas, Wayne Luk, Philip H.W. Leong, John D. Villasenor
November 2007 **ACM Computing Surveys (CSUR)**, Volume 39 Issue 4

Publisher: ACM

Full text available: pdf(426.75 KB) Additional Information: [full citation](#), [abstract](#),
[references](#), [index terms](#)

Rapid generation of high quality Gaussian random numbers is a key capability for simulations across a wide range of disciplines. Advances in computing have brought the power to conduct simulations with very large numbers of random numbers and with it, ...

Keywords: Gaussian, Random numbers, normal, simulation

13 Manufacturing supply chain applications 2: parameterization of fast and accurate simulations for complex supply networks

Brett Marc Duarte, John W. Fowler, Kraig Knutson, Esma Gel, Dan Shunk
December 2002 **WSC '02: Proceedings of the 34th conference on Winter simulation: exploring new frontiers**

Publisher: Winter Simulation Conference

Full text available: pdf(435.88 KB) Additional Information: [full citation](#), [abstract](#),
[references](#), [cited by](#)

More efficient and effective control of supply networks is conservatively worth billions of dollars to the world economy. Adopting an approach by which the basic disciplines of Industrial Engineering, Control Engineering, System Simulation and Business ...

14 Self-tuning cost modeling of user-defined functions in an object-relational DBMS



Zhen He, Byung Suk Lee, Robert Snapp
September 2005 **ACM Transactions on Database Systems (TODS)**,
Volume 30 Issue 3

Publisher: ACM

Full text available: pdf(1.42 MB) Additional Information: [full citation](#), [abstract](#), [references](#),
[cited by](#), [index terms](#)

Query optimizers in object-relational database management systems typically require users to provide the execution cost models of user-defined functions (UDFs). Despite this need, however, there has been little work done to provide such a model. The ...

Keywords: K-nearest neighbors, cost modeling, object relational DBMS, quadtree, query optimization, self-tuning

15 Wavelet synopses for general error metrics



Minos Garofalakis, Amit Kumar

December 2005 **ACM Transactions on Database Systems (TODS)**,

Volume 30 Issue 4

Publisher: ACM

Full text available: pdf(1.32 MB) Additional Information: [full citation](#), [abstract](#), [references](#),
[cited by](#), [index terms](#)

Several studies have demonstrated the effectiveness of the wavelet decomposition as a tool for reducing large amounts of data down to compact *wavelet synopses* that can be used to obtain fast, accurate approximate query answers. Conventional wavelet ...

Keywords: Data synopses, Haar wavelets, approximate query processing

16 XSKETCH synopses for XML data graphs



Neoklis Polyzotis, Minos Garofalakis

September 2006 **ACM Transactions on Database Systems (TODS)**,

Volume 31 Issue 3

Publisher: ACM

Full text available: pdf(885.57 KB) Additional Information: [full citation](#), [appendices and](#)
[supplements](#), [abstract](#),
[references](#), [index terms](#)

Effective support for XML query languages is becoming increasingly important with the emergence of new applications that access large volumes of XML data. All existing proposals for querying XML (e.g., XQuery) rely on a pattern-specification language ...

Keywords: XML, approximate query processing, data synopses, path expressions

17 Effective use of block-level sampling in statistics estimation



Surajit Chaudhuri, Gautam Das, Utkarsh Srivastava

June 2004 **SIGMOD '04: Proceedings of the 2004 ACM SIGMOD**

international conference on Management of data

Publisher: ACM

Full text available: pdf(415.81 KB) Additional Information: [full citation](#), [abstract](#),
[references](#)

Block-level sampling is far more efficient than true uniform-random sampling over a large database, but prone to significant errors if used to create database statistics. In this paper, we develop principled approaches to overcome this limitation of ...

18 Modeling skew in data streams

Flip Korn, S. Muthukrishnan, Yihua Wu



June 2006 **SIGMOD '06**: Proceedings of the 2006 ACM SIGMOD
international conference on Management of data

Publisher: ACM

Full text available: pdf(3.49 MB) Additional Information: [full citation](#), [abstract](#), [references](#),
[index terms](#)

Data stream applications have made use of statistical summaries to reason about the data using nonparametric tools such as histograms, heavy hitters, and join sizes. However, relatively little attention has been paid to *modeling* stream data parametrically, ...

Keywords: estimation, modeling, skew, streaming algorithms

19 NSL-BLRL: Efficient CacheWarmup for Sampled Processor

Simulation

Luk Van Ertvelde, Filip Hellebaut, Lieven Eeckhout, Koen De Bosschere

April 2006 **ANSS '06**: Proceedings of the 39th annual Symposium on
Simulation

Publisher: IEEE Computer Society

Full text available: pdf(263.33 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Architectural simulation is extremely time-consuming given the huge number of instructions that need to be simulated for contemporary benchmarks. Sampled simulation which selects a number of samples from the complete benchmark execution yields substantial ...

20 Space efficient mining of multigraph streams



Graham Cormode, S. Muthukrishnan

June 2005 **PODS '05**: Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems

Publisher: ACM

Full text available: pdf(340.94 KB) Additional Information: [full citation](#), [abstract](#),
[references](#), [cited by](#)

The challenge of monitoring massive amounts of data generated by communication networks has led to the interest in data stream processing. We study streams of edges in massive communication multigraphs, defined by (source, destination) pairs. The goal ...

Results 1 - 20 of 50

Result page: [1](#) [2](#) [3](#) [next](#) [>>](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2008 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)